

Discrete Quantum Control - State Preparation

Jon R. Grice and David A. Meyer

Department of Mathematics

*University of California, San Diego**

(Dated: April 28, 2012)

A discrete-time method for solving problems in optimal quantum control is presented. Controlling the time discretized Markovian dynamics of a quantum system can be reduced to a Markov-decision process. We demonstrate this method in this with a class of simple one qubit systems, which are also discretized in space. For the task of state preparation we solve the examples both numerically and analytically with dynamic programming techniques.

PACS numbers: 03.67.Ac, 02.30.Yy

I. INTRODUCTION

Quantum control is a discipline that will see its importance match any of the technologies that it enables – quantum information sciences, to take a prominent example.

But, the problems which arise in Quantum control can be very difficult to solve. This paper explores a possible direction towards formulating and solving problems from optimal quantum control by taking inspiration from quantum computing.

The control of quantum systems has been studied for decades, see e.g. [1]. Because real systems and measurements are often continuous [16], the usual technique in optimal quantum control is to study the stochastic master equation for the system being controlled. A sizable body of literature exists which details the control of continuous time quantum systems [6].

Some systems are inherently time-discrete like the Stern-Gerlach experiment or a quantum circuit, but any time-continuous system must be discretized if it is to be simulated on a quantum computer [8]. These time-discrete simulations are what we seek to control, in lieu of the continuous model of the system.

We model the system by slicing time into steps where the evolution on each step is given by a quantum circuit (including nonunitary quantum operations and measurements). Each timestep's circuit effects a Markovian step on the system's state. Since we want to control a Markov chain we turn to the theory of Markov decision processes (MDPs) [15], specifically to the field of control theory called Discrete-Time Optimal Stochastic Control [3].

A MDP consists of two parts, the agent and the system. The agent consists of the experimenter and a (classical) computer system, the controller. The system is an open quantum system or a model thereof, interacting through its environment, from which a measurement apparatus takes measurements which are reported to the agent. The agent in turn can vary the Hamiltonian (apply controls) which acts on the system to some specified degree. The closed loop process of taking measurements and applying controls is repeated for either a certain amount of time or indefinitely, depending on the goal the experimenter has in mind. The controller is programmed to use the information from the measurements along with the knowledge of the system dynamics to apply the appropriate controls in order to best further the goals of the experimenter. In this paper we discuss methods for finding the best program, on average, that the controller should use.

The MDPs that arise in this setting have a continuous state space and transition matrices which arise from the dynamics of the system which give them more structure than most MDPs that are studied. It is this structure that gives us more power to solve these systems, for example, we can solve some MDPs analytically.

We demonstrate solving a system by using two sample models, the affine-quadratic and the threshold models, so-named for the cost functions that determine their behavior. These models have a simple enough state space (the interval $[0, 1]$) that there is little difficulty in discretizing that space and solving the Bellman equations that arise numerically.

The discrete technique presented in this paper is also applicable to more complicated models. A subsequent paper will study open systems undergoing decoherence and bipartite systems undergoing disentanglement.

The two example models are chosen to be as elementary as possible, the system being that of a qubit undergoing a stochastic evolution in which the state remains pure. There is a desired target state that must be reached after a

*Electronic address: jgrice@math.ucsd.edu; Electronic address: dmeyer@math.ucsd.edu

fixed amount of time, but the target state is unstable.

We solve the models first numerically and note peculiar behavior of the solutions, such as the as-if piecewise linear shape of the decision function graph, and corresponding parabolic shape of the value functions, which do not necessarily have a minimum at the target state. This means it is not good to stay near the target state if it is unstable and the end of the experiment is far off. As the value functions are propagated backwards in time as is done in dynamic programming fast convergence is evident for the value functions tending to a steady-state solution. This means that for a very long experiment relative to the probability of decaying away from the target state, the controller takes little action until the latter part of the experiment.

Our parameterization of the measurement operators allows us to study the best measurements to make in order to control the system as efficiently as possible. Interestingly we find that having numerous (more than two) measurement operators can make the system easier to control, in the sense that the value functions are smaller compared to a system with only two measurement outcomes.

These behaviors of the affine-quadratic model that are observed from numerical simulations can be explained by the analytic solution of the Bellman equation. Due to the simplicity of the system the Bellman equation can be approximately solved with known error bounds.

A. The state preparation problem

Algorithms in quantum computation often require an initial state to be prepared. The state preparation module would fill this need by preparing the desired state $|\psi_T\rangle$ (as closely as possible) at the correct time by starting with an arbitrary unprepared state [10]. Oftentimes one can perform a projective measurement in which one of the outcomes is the desired state to be prepared. This however, may not happen often enough, fast enough [9], or there may not be the physical resources available to make the correct state with a single measurement.

In our two models the system begins in a known, but arbitrary pure state $|\psi_0\rangle$. After N timesteps, the final state $|\psi_N\rangle$ is evaluated against the target state $|\psi_T\rangle$ (there can be more than one target state in general). For the cost function formulation of the finite horizon MDP (see below), the success of the state preparation process of measured by the terminal cost function J_N^* . Clearly J_N^* should take its minima on the target states.

For the two models described in this paper, there is a single target state $|\psi_T\rangle = |0\rangle$. The first model we deal with, the affine model, we choose

$$J_N^{\text{aff}}(|\psi_N\rangle) = D_{\text{tr}}(|\psi_N\rangle, |\psi_T\rangle) = 1 - \langle\psi_N|\psi_T\rangle^2, \quad (1)$$

where D_{tr} is the trace distance. Suppose a quantum algorithm expects $|\psi_T\rangle$ as an input but gets $|\psi_N\rangle$ instead. Then that algorithm performs worst when J_N^{aff} is greatest, in the following sense. Suppose the quantum algorithm is represented by a unitary operation W applied to $|\psi_T\rangle$, followed by a POVM with operators $\{\Pi_i\}_i$. Outcome i is observed with probability $\text{tr} \Pi_i W^\dagger |\psi_T\rangle \langle\psi_T| W$. This provides the probability distribution $\{\text{tr} \Pi_i W^\dagger |\psi_T\rangle \langle\psi_T| W\}_i$. Taking the l^1 (classical trace) distance between that distribution and the one obtained by using $|\psi_N\rangle$ as the input state, $\{\text{tr} \Pi_i W^\dagger |\psi_N\rangle \langle\psi_N| W\}_i$, we can get a reasonable measure of success. Since we are not specifying the particular algorithm we do not know the POVM that will be applied. However the quantum trace distance $D_{\text{tr}}(|\psi_N\rangle, |\psi_T\rangle)$ provides an upper bound for the l^1 distance for any POVM [12]. Since we want the measure of success to be smaller when the probability distributions are similar, we get equation (1). Equation (1) varies with the overlap (squared fidelity) of $|\psi_N\rangle$ and $|\psi_T\rangle$, the probability of observing one when making a measurement looking for the other.

Open quantum systems can be approximated so that their environment “forgets” its interaction with the system very quickly (although in fact our simple model is closed). Thus Markov processes are used to model quantum systems and in our discrete case these will be Markov chains.

B. Markov decision processes

A Markov decision process is a controllable Markov process. We will present MDPs in a slightly different way than is usual, to make it easier to adapt the quantum model to the formalism of the MDP. All Markov processes in this paper will be Markov chains, such that the time set $I = \{0, \dots, N-1\}$ is finite. A MDP is a set of states \mathcal{S} together with sets of actions $A(s, i) \subseteq \mathcal{S}$ that can be taken at each state $s \in \mathcal{S}$ and time $i \in I$ and transition probabilities $P_{\hat{s}, s'}$. If the current state at time i is s then $P_{\hat{s}, s'}$ is the probability that state $\hat{s} \in A(s, i)$ transitions to state s' in the next timestep where the action \hat{s} is an intermediary state that is chosen by the agent out of the set $A(s, i)$. The agent chooses the action by following a policy, a function $\pi : \mathcal{S} \times I \rightarrow \mathcal{S}$ such that $\pi(s, i) \in A(s, i)$. (There are formalizations of the policy function which enable the agent to act probabilistically and in a non-Markovian way but this level of generality is not needed in this paper [17].)

The model allows only one way to take the state s to \hat{s} at time i and it is assigned a control cost $g_i(s, \hat{s})$. An alternate formulation of MDPs uses rewards (negative of costs) associated to the state resulting at the next timestep. If $g_i(s, \hat{s})$ is independent of time the subscript i is dropped.

In summary, a single timestep proceeds as follows:

$$s_i \xrightarrow{\pi_i} \hat{s}_i \xrightarrow{\omega_i} s_{i+1} \quad (2)$$

where ω_i is a random outcome choosing $s'_{i+1} = s_{i+1}$ with probability $P_{\hat{s}_i, s'_{i+1}}$.

Markov decision processes can be run for a definite or indefinite amount of time. With the state preparation problem, the experiment runs for a fixed amount of time (N timesteps). This is called the finite-horizon problem. Experiments which run for an indefinite amount of time, which can stop at a time unknown to the experimenter or run for an infinite amount of time ($I = \mathbb{N}$) are interesting and arise in Quantum control but the theory of modeling such systems with MDPs will not be discussed in this paper.

For a finite-horizon problem, the state of the system at the end of the experiment s_N is of importance. The terminal cost function $J_N^*(s_N)$ evaluates s_N . If the experiment starts with initial state s_0 , with transition outcomes ω and the controller follows the policy π then the total cost is

$$G(s_0, \pi, \omega) = \sum_{i=0}^{N-1} g(s_i, \hat{s}_i) + J_N^*(s_N). \quad (3)$$

where the s_i are obtained from the procedure (2). Let

$$J_0^\pi(s_0) = \mathbb{E}[G(s_0, \pi, \omega)], \quad (4)$$

the expected total cost of the experiment, following policy π with initial state s_0 . For a fixed s_0 we would like to find an optimal policy π^* that minimizes $J_0^\pi(s_0)$. In that case we write $J_0^*(s_0) = J_0^{\pi^*}(s_0)$.

The Dynamic programming algorithm [2, 3, 15] is a well-known technique for finding the optimal policies π^* . To state it, we define the optimal cost-to-go at time k ,

$$J_k^*(s) = \mathbb{E} \left[\sum_{i=k}^{N-1} g(s_i, \hat{s}_i) + J_N^*(s_N) \right], \quad (5)$$

where, as before, the s_i are obtained from the process (2). The optimal cost-to-go tells us the minimum remaining total cost when starting in state s at time i . Under suitable conditions [3],

$$J_i^*(s) = \mathcal{U}[J_{i+1}^*(s)] = \inf_{\hat{s} \in A(s, i)} \left\{ \mathbb{E} [g(s, \hat{s}) + J_{i+1}^*(s')] \right\}, \quad (6)$$

for $0 \leq i < N$. The dynamic programming algorithm (6) states that the optimal cost-to-go functions can be computed by iterating backwards in time.

We will be concerned with cases when the infimum in equation (6) is achieved and so (an) optimal policy $\pi^*(s, i) = \hat{s}$ is obtained from taking an argmin in equation (6). In the next section we will apply this framework to a general discretized physical system.

C. The state preparation problem as a Markov decision process

By truncating the Baker-Campbell-Hausdorff (BCH) formula or using the Trotter formula [5, 8, 12, 14] on the system dynamics of a finite-dimensional open quantum system, a discrete time approximation can be obtained for the system:

$$s_i \longrightarrow \boxed{\text{Controllable } U_c} \longrightarrow \boxed{\text{Free } U_f} \longrightarrow \boxed{\text{Meas.}} \longrightarrow \boxed{\text{Noise}} \longrightarrow s_{i+1}. \quad (7)$$

The circuit (7) represents a single timeslice of system evolution. It is applied repeatedly to the starting state s_0 . The ordering of the blocks in (7) becomes less relevant as the length of the timeslice $\Delta t \rightarrow 0$. Questions arise pertaining to the behavior of the discrete model in the continuous time limit. These issues are discussed in [5].

The structure of the circuit mirrors our formulation of a MDP. The state s_i enters from the left. The U_c block is the part where the agent takes an action, choosing a new state $\hat{s}_i \in A(s, i)$. The action is the state resulting from a

unitary operation applied to s chosen from a set of control Hamiltonians \mathcal{H}_i that the agent has available at time i . Specifically, $A(s, i) = \{e^{iH\Delta t}s \mid H \in \mathcal{H}_i\}$.

If $A(s, i)$ does not depend on s and i then we can assume that any desired action is in $A(s, i)$, and disallow physically unrealistic or undesirable actions at a given state by making the cost prohibitive. In that case (and in the examples) we effectively set $A(s, i) = \mathcal{S}$. If the action is \hat{s} , the cost paid at timestep i is $g(s, \hat{s})$. The cost function will be explained in a later section. If there are multiple Hamiltonians in \mathcal{H}_i which result in the same action we choose the best one under the same physical constraints that the cost function was chosen (in order to make the cost function well-defined).

The free evolution unitary U_f is the next step. This is coherent evolution of the system that cannot be controlled directly by the controller. This step can be folded into the previous step by composing U_f with the action.

The next two steps, measurement and noise can also be composed together into a single non-trace preserving quantum operation that sends the state \hat{s} to one of the states $\{R_\alpha(\hat{s})\}_{\alpha \in \Omega} \subseteq \mathcal{S}$ with probability distribution $\Pr(\alpha|\hat{s}) = P_{\hat{s}, R_\alpha(s')}$ and set of measurement outcomes Ω . In the finite dimensional system we consider, Ω will be a finite set. Then the Dynamic programming algorithm (6) becomes

$$J_i^*(s) = \min_{\hat{s} \in A(s, i)} \left\{ \int_{\Omega} g(s, \hat{s}) + J_{i+1}^*(R_\alpha(\hat{s})) \, d\Pr(\alpha|\hat{s}) \right\}, \quad (8)$$

A few properties of the Bellman iteration $\mathcal{U}[\cdot]$ from (6) that are useful for analyzing the examples will be summarized.

Proposition 1. *Let \mathcal{S} be compact and Ω be finite. Then $\mathcal{U}[\cdot]$ is a map on $\mathcal{C}(\mathcal{S})$. That is, $\mathcal{U}[J]$ is continuous if J is. Additionally,*

- I. *If $g(s, \hat{s}) \geq 0$ for all $s, \hat{s} \in \mathcal{S}$ with equality for $s = \hat{s}$ then the constant function $J(s) \equiv C$ is a fixed point for $\mathcal{U}[\cdot]$ for all C .*
- II. *$\|\mathcal{U}[J] - \mathcal{U}[V]\|_\infty \leq \|J - V\|_\infty$ In particular, $\mathcal{U}[\cdot]$ is a continuous map.*

Proof. Item II: Fix s . Then

$$\begin{aligned} |\mathcal{U}[J](s) - \mathcal{U}[V](s)| &= \left| \min_{\hat{s}} \{g(s, \hat{s}) + \sum J(R_\alpha(\hat{s})) \Pr(\alpha|\hat{s})\} \right. \\ &\quad \left. - \min_{\hat{s}} \{g(s, \hat{s}) + \sum V(R_\alpha(\hat{s})) \Pr(\alpha|\hat{s})\} \right| \\ &\leq \max_{\hat{s}} \left| \sum (J(R_\alpha(\hat{s})) - V(R_\alpha(\hat{s}))) \Pr(\alpha|\hat{s}) \right| \\ &\leq \max_z |J(z) - V(z)| \sum \Pr(\alpha|\hat{s}) \\ &= \|J - V\|_\infty. \end{aligned}$$

□

We expect proposition 1 is well-known in the field of discrete-time stochastic control but are unable to find sources. We now turn our attention to the specific models.

II. OPTIMAL CONTROL PROBLEMS FOR CYLINDRICALLY SYMMETRIC QUBIT STATES

We will demonstrate this technique with two single qubit models. In order to make the optimization procedure needed to solve the Bellman equation straightforward, we choose one of the simplest one-qubit systems, one that can be parameterized by a single real number. One way to satisfy this constraint is to confine an arbitrary state s in the system \mathcal{S} to half of a great circle on the surface of the Bloch sphere \mathcal{B} . We choose a basis $\{|0\rangle, |1\rangle\}$ so that if $s_x \in \mathcal{S}$, then

$$s_x = \sqrt{1-x}|0\rangle + \sqrt{x}|1\rangle, \quad (9)$$

for $x \in [0, 1]$. We will hereafter abuse the notation by identifying the system \mathcal{S} with its parameterization $[0, 1]$ and state s_x with its parameter x . When an additional parameter ϕ is introduced to represent relative phase the entire Bloch sphere is swept out by rotating the set \mathcal{S} about the z -axis (the axis containing $|0\rangle$ and $|1\rangle$) as ϕ is varied. We may therefore consider any state $s \in \mathcal{S}$ to be the equivalence class $\{\mathcal{R}_z(\phi)s\}$ (\mathcal{R}_z is the rotation about the z -axis operator) making the set of equivalence classes \mathcal{S} a cylindrically symmetric version of \mathcal{B} .

There are several reasons why our cylindrically-symmetric \mathcal{S} is not an unrealistic choice. If the experiment requires the preparation of a unpopulated state $|0\rangle$ without any consideration to the relative phase, then the terminal cost function will be cylindrically symmetric. If the control step is strobing an appropriately aligned magnetic field, (for example, one with no σ_x or σ_z in the Hamiltonian) then the control cost function g is cylindrically symmetric. In fact, there are several families of physically-motivated operators which preserve \mathcal{S} and hence are cylindrically symmetric.

A. Operators on \mathcal{S}

We can think of a noisy quantum operation $\mathcal{E} : s \mapsto s'$ being applied to a system $\rho_s = |\psi_s\rangle\langle\psi_s|$ as the system interacting with its environment ρ_e . This can be modeled by an entangling unitary operator U_{se} acting on the environment-system ρ_{se} followed by a measurement \mathcal{M}_e on the environment. When the record of \mathcal{M}_e is ignored the resultant system state $\rho_{s'} = \text{tr}_e(\sum_j M_j U_{se} \rho_s U_{se}^\dagger M_j^\dagger)$ is left as a mixed state meaning that in general $s' \notin \mathcal{S}$. Now suppose the experimenter is allowed to monitor the environment (for example, the environment is an electromagnetic channel monitored with a photodetector [13]) and know which M_j occurred. Because in that case $s' \in \mathcal{S}$ we will use the latter scenario in our model to derive some physically motivated operations on \mathcal{S} . As explained below some important quantum operations (amplitude damping, phase damping, etc.) have Kraus operators which preserve \mathcal{S} and these operators can serve as measurement operators for our system.

Considering the set of states as column vectors, $\{|\psi_{s_x}\rangle = \begin{pmatrix} \sqrt{1-x} \\ \sqrt{x} \end{pmatrix} \mid s_x \in \mathcal{S}\}$ is preserved by the real nonzero nonnegative-entry 2×2 matrices $M_2(\mathbb{R}^+) \setminus \{0\}$, up to a normalization factor.

A finite subset $\{K_j\} \subset M_2(\mathbb{R}^+)$ satisfying the completeness relation

$$\sum_j K_j^\dagger K_j = I_2, \quad (10)$$

can represent the operators for either a complete measurement or a trace-preserving quantum operation [12] (we think of it as the former). It is then a simple calculation to check that the K_j belong to either or both the filtering class (rank 2)

$$\mathcal{F} = \left\{ \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}, \begin{pmatrix} 0 & a \\ b & 0 \end{pmatrix} \mid ab \neq 0 \right\}$$

or the jump class (rank 1)

$$\mathcal{J} = \left\{ \begin{pmatrix} 0 & a \\ 0 & b \end{pmatrix}, \begin{pmatrix} a & 0 \\ b & 0 \end{pmatrix} \right\}.$$

The operators in \mathcal{F} arise from the Bayesian updating of the state conditioned on the result of the measurement while the operators in \mathcal{J} are direct jumps (e.g. wave function collapse) to a fixed state.

An operator $K \in \mathcal{F} \cup \mathcal{J}$ is also considered a map on \mathcal{S} ,

$$K : \hat{s} \mapsto s'$$

such that

$$|\psi_{s'}\rangle = \frac{K|\psi_{\hat{s}}\rangle}{\|K|\psi_{\hat{s}}\rangle\|}.$$

Furthermore if K is taken to be a measurement operator then $\text{Pr}(K|\hat{s}) = \|K|\psi_{\hat{s}}\rangle\|^2$ for all $\hat{s} \in \mathcal{S}$. Table I describes the measurement operators in that sense.

Since the discrete systems in this paper are approximations of real continuous systems, in a future paper we investigate such operators in the infinitesimal time limit—the filtering operators go smoothly to the identity operator while the jump operators remain fixed jumps (although the probabilities of occurrence go to zero).

Some examples are [12] amplitude damping $\{f_1(1, 1-\gamma), j_1(\gamma, 0)\}$, phase damping $\{f_1(1, 1-\gamma), j_1(0, \gamma)\}$, generalized amplitude damping $\{f_1(p, p(1-\gamma)), f_1((1-p)(1-\gamma)), j_1(p\gamma, 0), j_2(0, (1-p)\gamma)\}$ and the bit-flip channel $\{f_1(p, p), f_2(1-p, 1-p)\}$.

B. Filtering and Jump operators

When $M_i \in \mathcal{J}$ for all $0 \leq i < m$, the outcomes are in m discrete states, call them $|\psi_i\rangle$, $0 \leq i < m$, possibly nondistinct. There may be interesting problems when the set of the $|\psi_i\rangle$ are nonorthogonal, but we will not consider jump-only measurements further in this paper.

TABLE I: Measurement operators

Name	Operator K	s'	$\Pr(K s)$
$f_1(a, b)$	$\begin{pmatrix} \sqrt{a} & 0 \\ 0 & \sqrt{b} \end{pmatrix}$	$\frac{bs}{(b-a)s+a}$	$(b-a)s+a$
$f_2(a, b)$	$\begin{pmatrix} 0 & \sqrt{b} \\ \sqrt{a} & 0 \end{pmatrix}$	$\frac{a(1-s)}{(b-a)s+a}$	$(b-a)s+a$
$j_1(a, b)$	$\begin{pmatrix} 0 & \sqrt{a} \\ 0 & \sqrt{b} \end{pmatrix}$	$\frac{b}{a+b}$	$(a+b)s$
$j_2(a, b)$	$\begin{pmatrix} \sqrt{a} & 0 \\ \sqrt{b} & 0 \end{pmatrix}$	$\frac{b}{a+b}$	$(a+b)(1-s)$

The measurement sets \mathcal{M} that we will use will contain at least one filtering operator and zero or more jump measurements so that \mathcal{M} contains two or more measurement operators. Experience suggests that having more than two nonorthogonal measurement operators can make the system easier to control.

III. COST FUNCTIONS FOR A SIMPLE MODEL

The two examples that we discuss in this paper differ only by their cost functions. For the first, the affine model, the terminal cost function is equation (1), which, when recast in the x -parameterization is

$$J_N^{\text{aff}}(x_N) = |x - x_T| = x, \quad (11)$$

where the target state $x_T = 0$ if $s_T = |0\rangle$. Other functions which have their global minima at the target states could be justified. One might decide that a resultant state more than 0.2 units from the target state in \mathcal{S} is unacceptable. Then we might have a J_N^* like

$$J_N^{\text{thresh}}(x) = \begin{cases} |x - x_T|, & \text{if } |x - x_T| \leq 0.2 \\ 100, & \text{if } |x - x_T| > 0.2. \end{cases} \quad (12)$$

The number 100 was chosen to be much larger than typical values of the cost-to-go functions (essentially infinity) but not large enough to cause overflows in computation.

With our discrete time formulation of the control problem, we do not directly deal with differential equations and so nonsmooth functions like J_N^{thresh} can be incorporated naturally into the model.

A. Cost-to-move function

In addition to the terminal cost there are resources that are used by the controller at every timestep to drive the system. The use of these resources is quantified by the cost-to-move function. For simplicity, as mentioned in section IC, the cost-to-move function $g(s, \hat{s})$ will only depend in the current state s and the state \hat{s} that the controller will move the system to in the next timestep by using one of its Hamiltonians.

Part of what makes quantum control hard is that the controller can not move the system to any state that it chooses in a single step. Otherwise, the step preparation problem would be trivial. The transitions $s \mapsto \hat{s}$ which are inadmissible (disallowed) in the model are due to physical considerations. The Hamiltonians that the controller has available, the size of the timestep and the way the system is discretized determine which pairs (s, \hat{s}) are inadmissible. In that case, we assign

$$g(s, \hat{s}) = \infty$$

for an inadmissible pair (s, \hat{s}) . As in equation (12), for computational purposes ∞ is taken to be a large number. The function $J_i^{\hat{s},*}$ is the cost-to-go if the next action is \hat{s} and then an optimal policy is followed thereafter. Optimizing $J_i^{\hat{s},*}$ over \hat{s} is the same as the optimization step in the dynamic programming algorithm (DPA) equation (6) and so it often shows up in implementations of the DPA. The advantage in making all pairs in $\mathcal{S} \times \mathcal{S}$ admissible and redefining infinity to a large number is that it makes the domain of $J_i^{y,*}(x)$ (in variables (x, y)) rectangular and finite on the entire domain, easing the optimization process.

As for reasonable cost functions, the most straightforward metric, for states on the Bloch sphere \mathcal{B} , might be the angular distance. Two states $x, y \in \mathcal{S}$ have angular distance

$$d_\theta(x, y) = \arccos(1 - 2y - 2x + 4xy + 4\sqrt{x(1-y)y(1-y)})$$

One might imagine that in the U_c control step above the controller has some Hamiltonians that can be switched on and off for a percentage of the time afforded to the controller. Then perhaps all states that are a certain distance away in the d_θ metric can be reached, for a similar use of resources. We might have then

$$g^{\text{thresh}}(x, y) = \begin{cases} 0, & \text{if } d_\theta(x, y) < 0.2 \\ 100, & \text{if } d_\theta(x, y) \geq 0.2. \end{cases} \quad (13)$$

On a similar vein to the terminal cost, 100 is an arbitrary ‘big’ number so that if $J_0^*(x) \geq 100$ then there is no way to control the system from x to x_T in the model.

Other possible cost functions are d_θ or d_θ^2 . Here we imagine that applying the Hamiltonians in the control step U_c cost energy which is proportional to how far in the d_θ metric that the state has moved [11]. In this case we would like to balance the usage of energy to the closeness of the resultant state at the end of the experiment to the target state. In this situation, having infinite energy ($g(s, \hat{s}) = 0$) might allow the controller to make very strong pulses toward the end of the experiment. Whether or not that is an optimal control solution depends on the probability of states near the target state decaying away before the end of the experiment. In [4] the authors faced a similar problem.

In order to make the system solvable analytically, we will use the second order Taylor polynomial of d_θ^2 expanded about $(\frac{1}{2}, \frac{1}{2})$, giving a quadratic symmetric in its arguments. It is

$$g^{\text{aff}}(x, y) = 4(x - y)^2. \quad (14)$$

Using the function $C(s, \hat{s}) = |s - \hat{s}|$ would be even simpler but leads to trivial controls when equation (6) is solved for the example model in the next section.

IV. SOLVING THE DPA FOR THE MODELS

Putting the affine terminal cost (11) and the quadratic movement cost (14) into equation (6) we get

$$J_{i-1}^*(x) = \min_{0 \leq y \leq 1} \left\{ 4(x - y)^2 + \sum_j J_i^* \left(\frac{a_j y + b_j}{c_j y + d_j} \right) (c_j y + d_j) \right\}, \quad (15)$$

$$J_N^*(x) = x,$$

with constants a_j , b_j , c_j and d_j which depend on the measurement set chosen from table I. We will also solve equation (6) numerically using the threshold cost functions (12) and (13).

In order to solve equations (15) numerically, we must specify the measurement set to be used. Numerous examples that we have experimented with have the same basic behavior, which is typified by the system with measurement operators $\{f_1(a, 1), j_2(0, 1 - a)\}$. With some rescaling equation (15) becomes

$$\mathcal{U}[J](x) = \min_{0 \leq y \leq 1} \left\{ (x - y)^2 + J \left(\frac{y}{(1 - a)y + a} \right) \left((1 - a)y + a \right) + J(1) \left(1 - ((1 - a)y + a) \right) \right\}. \quad (16)$$

The parameter $0 \leq a \leq 1$ determines how unstable the target state is. If the system is a two-level atom then it is the rate of decay. We can vary a in order to study how the optimal policies change with the instability of the system.

A. Solving numerically

A simple way to proceed numerically is to discretize \mathcal{S} into say, 100 equally spaced points. For the resolution desired this naive assumption turns out to be acceptable for reasons that will become clear in the analytic solution section. The functions J_i^* can be approximated as a discrete list of values and computed recursively by minimizing over a discrete list. Figure IV Aa. shows a graph of J_i^* for $i = 0, 1, \dots, 5$ when $a = 0.8$. Similar graphs are obtained for $0 < a < 1$ and for systems with other sets of measurement operators. Figure IV Ac. shows the optimal policies $\pi^*(x, i)$ for the affine problem. From the picture, the $\pi^*(x, i)$ appear to be piecewise affine. The analytic solution will tell us that it is not the case, although the $\pi^*(x, i)$ are closely approximated by piecewise affine functions.

When we use the threshold functions, a typical run might look like figure IV A(b,d.). The cost-to-go functions are quite different for the two models. One can easily see how the cutoff in J_N^* thresh propagates through the J_i^* . We expect an analytic analysis of this model to be fairly easy. Another interesting feature which is shared with the latter model is that the local minima for the various value functions do not lie at the target state. This is due to the large probability of a state near the target state decaying away before the end of the experiment.

B. Approximating the solution analytically

We can explain some of the behaviors that arose with the affine/quadratic model (16). The goal is to show that if cost-to-go J is quadratic then $\mathcal{U}[J]$ is approximately quadratic. Appealing to the continuity of the iteration map, we can obtain reasonable approximations to $\mathcal{U}^j[J]$ for small values of j . Due to constant-function fixed points of $\mathcal{U}[\cdot]$, for larger values of j , $\mathcal{U}^j[J]$ becomes uninteresting.

Suppose $J(x) = Ax^2 + Bx + C$. Considering the sum in equation (15), we get linear terms from jump-type measurements and for each filtering-type measurement operator we get a term like $J(\frac{ay+b}{cy+d})(cy+d)$. Evaluating this, we get a quadratic term and a fractional term $N(y)$ given by

$$N(y) = \frac{A(ad - bc)^2}{c^2(d + cy)}. \quad (17)$$

The second order Lagrange interpolating polynomial agreeing with $N(y)$ at the points $y_i = 0, \frac{1}{2}, 1$ is

$$L(y) = \frac{A(ad - bc)^2}{c^2} \left(\frac{1}{d} - \frac{c(3c + 2d)}{d(c + d)(c + 2d)}y + \frac{2c^2}{d(c + d)(c + 2d)}y^2 \right)$$

with error bounds [7]

$$|N(y) - L(y)| \leq \frac{\max_{0 \leq \xi \leq 1} |N^{(3)}(\xi)|}{3!} |y(y - \frac{1}{2})(y - 1)|.$$

The maximum error is bounded by $\left| \frac{\alpha}{(d+c)^4 12\sqrt{3}} \right|$. For an operator $f_1(a, b)$ or $f_2(a, b)$, the bound becomes

$$\left| \frac{Aa^2(b - a)}{b^2 12\sqrt{3}} \right|.$$

For example, if J was an affine function ($A = 0$) then the approximation is exact.

Let the approximate to $\mathcal{U}[J]$ be called $\tilde{\mathcal{U}}[J]$. Computing $\tilde{\mathcal{U}}[J](x)$ amounts to finding the minimum of a quadratic parameterized by x on a closed interval. $\tilde{\mathcal{U}}[J]$ is thus a piecewise quadratic function. Likewise to approximate $\mathcal{U}^2[J]$ we compute $\tilde{\mathcal{U}}[\tilde{\mathcal{U}}[J]]$ and get a bound on the error. We do not find the occasion to repeat this procedure to find, for example an approximation to J_0^* due to the algebra involved and the efficacy of the numerical method. The analytic solution explains why the graphs of the π^* look piecewise linear – because the graphs of the J_i^* look piecewise quadratic.

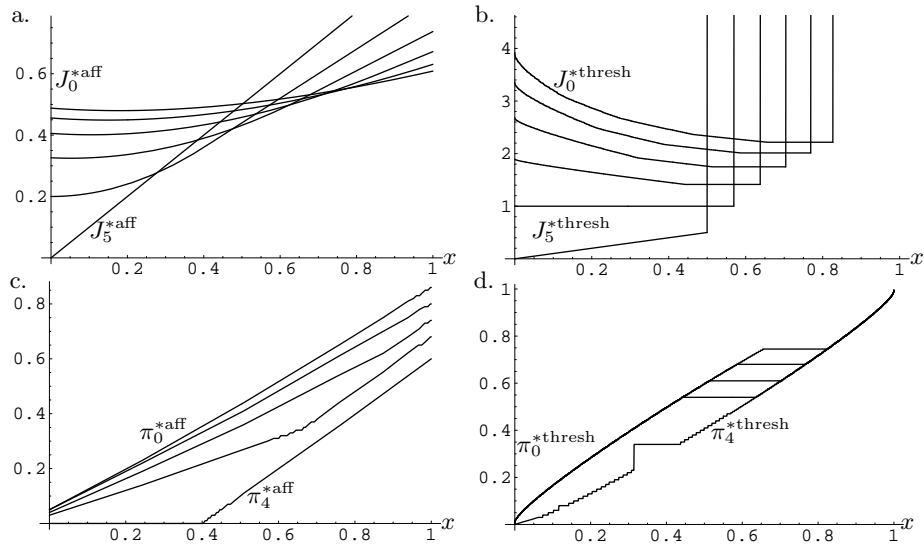


FIG. 1: Example costs-to-go and optimal policies.

V. CONCLUSION

We have demonstrated a method for controlling a quantum system by controlling its simulation. We have shown that the method works by applying it to two simple models. The method has been applied to more complicated models: a qubit undergoing decoherence, which has a more complicated state space, and two qubits undergoing disentanglement. These situations will appear in a future paper.

We need to resolve issues with the discretization. How small does the timeslice have to be to get ‘good’ control? How does the MDP turn into a stochastic differential equation as the timeslice goes to 0?

We are concerned with how the control problem scales with the size of the system. When searching the state space in the optimization step we rely on good parameterizations of the system. These become harder to come by as the size of the system increases.

Acknowledgments

This work has been partially supported by the National Science Foundation under grant ECS-0202087 and by the Defense Advanced Research Projects Agency as part of the Quantum Entanglement Science and Technology program under grant N66001-09-1-2025.

-
- [1] V.P. Belavkin. Theory of the control of observable quantum systems. *Autom. remote control*, 44:178–188, 1983.
 - [2] Richard Bellman. *Dynamic Programming*. Princeton University Press, 1957.
 - [3] D. Bertsekas and S. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific, 1996.
 - [4] L. Bouten, S. Edwards, and V. Belavkin. Bellman equations for optimal feedback control of qubit states. *J. Phys. B: At. Mol. Opt. Phys.*, 38:151–160, 2005.
 - [5] L. Bouten and R. Van Handel. Discrete approximation of quantum stochastic models. *eprint: arXiv:0803.4383v1 [math-ph]*, 2008.
 - [6] L. Bouten, R. van Handel, and M. James. A discrete invitation to quantum filtering and feedback control. *arXiv:math/0606118v4*, 2006.
 - [7] Burden and Faires. *Numerical Analysis*. Thomson Brooks/Cole, 8th edition, 1995.
 - [8] R. Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21:467–488, 1982.
 - [9] Kurt Jacobs. How to project qubits faster using quantum feedback. *Phys. Rev. A*, 67:030301, 2003.
 - [10] Kurt Jacobs. Applications of feedback control in quantum systems. *Proceedings of the 6th Asian Control Conference*, page 35, 2006.
 - [11] N. Margolus and L. Levitin. The maximum speed of dynamical evolution. *Physica D*, 120:188–195, 1998.
 - [12] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
 - [13] Saleh and Teich. *Fundamentals of Photonics*. Wiley-Interscience, 1991.
 - [14] A.T. Sornborger and E.D. Stewart. Higher order methods for simulations on quantum computers. *Phys.Rev. A*, 60:1956, 1999.
 - [15] R. Sutton and A. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
 - [16] Ramon van Handel, John K. Stockton, and Hideo Mabuchi. Modeling and feedback control design for quantum state preparation. *To appear: J. Opt. B (Special Issue: Quantum Control)*, 2005.
 - [17] J. van der Wal. *Stochastic dynamic programming : successive approximations and nearly optimal strategies for Markov decision processes and Markov games*. Amsterdam : Mathematisch Centrum, 1981.